

FORM PTO-1390
(REV. 11-94)U.S. DEPARTMENT OF COMMERCE
PATENT AND TRADEMARK OFFICE

ATTORNEY'S DOCKET NUMBER

**TRANSMITTAL LETTER TO THE UNITED STATES
DESIGNATED/ELECTED OFFICE (DO/EO/US)**

9392-023-999

09/890185INTERNATIONAL APPLICATION NO.
PCT/US00/02011INTERNATIONAL FILING DATE
January 28, 2000PRIORITY DATE CLAIMED
January 28, 1999

TITLE OF INVENTION

A METHOD AND SYSTEM FOR ROUTING CONTROL IN COMMUNICATION NETWORKS AND FOR SYSTEM CONTROL

APPLICANT(S) FOR DO/EO/US

Stuart A. KAUFFMAN and Tony A. PLATE

Applicant herewith submits to the United States Designated/ Elected Office (DO/EO/US) the following items under 35 U.S.C. 371:

1. This is a **FIRST** submission of items concerning a filing under 35 U.S.C. 371.
2. This is a **SECOND** or **SUBSEQUENT** submission of items concerning a filing under 35 U.S.C. 371.
3. This is an express request to begin national examination procedures (35 U.S.C. 371(f)) at any time rather than delay examination until the expiration of the applicable time limit set in 35 U.S.C. 371(b) and PCT Articles 22 and 39(1).
4. A proper Demand for International Preliminary Examination was made by the 19th month from the earliest claimed priority date.
5. A copy of the International Application as filed (35 U.S.C. 371(c)(2))
 - a. is transmitted herewith (required only if not transmitted by the International Bureau).
 - b. has been transmitted by the International Bureau.
 - c. is not required, as the application was filed in the United States Receiving Office (RO/US)
6. A translation of the International Application into English (35 U.S.C. 371(c)(2)).
7. Amendments to the claims of the International Application under PCT Article 19 (35 U.S.C. 371(c)(3))
 - a. are transmitted herewith (required only if not transmitted by the International Bureau).
 - b. have been transmitted by the International Bureaus.
 - c. have not been made; however, the time limit for making such amendments has NOT expired.
 - d. have not been made and will not be made.
8. A translation of the amendments to the claims under PCT Article 19 (35 U.S.C. 371(c)(3)).
9. An oath or declaration of the inventor(s) (35 U.S.C. 371(c)(4)).
10. A translation of the annexes to the International Preliminary Examination Report under PCT Article 36 (35 U.S.C. 371(c)(5)).

Items 11. to 16. below concern document(s) or information included:

11. An Information Disclosure Statement under 37 CFR 1.97 and 1.98.
12. An assignment document for recording. A separate cover sheet in compliance with 37 CFR 3.28 and 3.31 is included.
13. A **FIRST** preliminary amendment.
 A **SECOND** or **SUBSEQUENT** preliminary amendment.
14. A substitute specification.
15. A change of power of attorney and/or address letter.
16. Other items or information:

Power of Attorney by Assignee

INTERNATIONAL APPLICATION NO.
PCT/US00/02011

09/890185

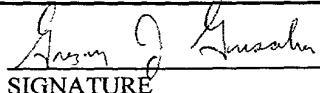
INTERNATIONAL FILING DATE
January 28, 200017. The U.S. National Fee (35 U.S.C. 371(c)(1)) and other fees as follows:

CLAIMS						
(1)FOR	(2)NUMBER FILED	(3)NUMBER EXTRA	(4)RATE	(5)CALCULATIONS		
TOTAL CLAIMS	50 - 20	30	X \$ 18.00	\$ 540.00		
INDEPENDENT CLAIMS	7 - 3	4	X \$ 80.00	320.00		
MULTIPLE DEPENDENT CLAIM(S) (if applicable)		+ \$ 270.00	<input type="checkbox"/>			
BASIC NATIONAL FEE (37 CFR 1.492(a)(1)-(5)):						
CHECK ONE BOX ONLY						
<input type="checkbox"/> International preliminary examination fee paid to USPTO (37 CFR 1.482) \$ 690						
<input type="checkbox"/> No international preliminary examination fee paid to USPTO (37 CFR 1.482) but international search fee paid to USPTO (37 CFR 1.445(a)(2)) \$ 710						
<input type="checkbox"/> Neither international preliminary examination fee (37 CFR 1.482) nor international search fee (37 CFR 1.445(a)(2)) paid to USPTO \$ 1000						
<input checked="" type="checkbox"/> International preliminary examination fee paid to USPTO (37 CFR 1.482) and all claims satisfied provisions of PCT Article 33(2) to (4) \$ 100						
<input type="checkbox"/> Filing with EPO or JPO search report \$ 860						
Surcharge of \$130.00 for furnishing the National fee or oath or declaration later than 20 30 mos. from the earliest claimed priority date (37 CFR 1.492(e)).						
				TOTAL OF ABOVE CALCULATIONS = 960.00		
Reduction by 1/2 for filing by small entity, if applicable. Affidavit must be filed also. (Note 37 CFR 1.9, 1.27, 1.28).						
				SUBTOTAL = 960.00		
Processing fee of \$130.00 for furnishing the English Translation later than 20 30 mos. from the earliest claimed priority date (37 CFR 1.492(f)).						
				TOTAL FEES ENCLOSED \$ 960.00		

a. A check in the amount of \$ to cover the above fees is enclosed.

b. Please charge Deposit Account No. 16-1150 in the amount of \$960.00 to cover the above fees. A copy of this sheet is enclosed.

c. The Commissioner is hereby authorized to charge any additional fees which may be required, or credit any overpayment to Deposit Account No. 16-1150. A copy of this sheet is enclosed.

18. Other instructions
n/a19. All correspondence for this application should be mailed to
PENNIE & EDMONDS LLP
1667 K STREET, N.W.
WASHINGTON, D.C. 2000620. All telephone inquiries should be made to (202) 496-4720Gregory J. Gonsalves
for Francis E. Morris
NAME
SIGNATURE

43,369

24,615

REGISTRATION NUMBER

July 30, 2001

DATE

A METHOD AND SYSTEM FOR ROUTING CONTROL IN
COMMUNICATION NETWORKS AND FOR SYSTEM CONTROL

FIELD OF THE INVENTION

The present invention relates generally to a method and system for routing control in communication networks and for system control. More particularly, the present invention performs routing by controlling the components in a network with software agents operating in a reward framework using p , tau, and patches to improve communication performance.

10

Background

Modern data-communication networks, as a non-limiting example packet-switched data networks, often present many potential routes between nodes that wish to communicate. Decisions about the route that data should take are usually made in a decentralized fashion by routers at the nodes. Decisions must be decentralized both because a centralized routing device would make the network vulnerable to single-point failures and because it would be impractical to communicate routing decisions from a centralized device to all the nodes in a spatially disperse network. Ideally, routing decisions should take into account both network topology (e.g., finding the shortest or least-cost path between two nodes) and current and historical network load (i.e., finding paths that do not utilize currently or historically overloaded communication links).

However, it is difficult to construct routers that make effective decisions based on load due to the problem of oscillation. For example, if link A is currently overloaded and link B is currently under loaded, then link B appears preferable to all the routers, which leads to link B being overloaded and link A being under loaded, and so on. Consequently, currently-fielded commercially-available routers take into account only network topology when making

routing decisions (though they may try to split traffic among equal-cost paths.) As a result, communication performance is not as good as is theoretically possible. Bandwidth, delay (latency) and reliability (i.e., packet loss) are all 5 negatively affected by routing decisions that do not take network load into account.

Accordingly, there is a pressing need for decentralized routing algorithms that can effectively take both network topology and current and historical load on communication links into account.

10

Summary of the Invention

The present invention present a method and system for routing control in communication networks by controlling the components in a network with software agents operating in a reward framework using p , τ , and patches to improve 15 communication performance.

The present invention includes a method for routing packets of data through a network of a plurality of components comprising the steps of:

controlling one or more of said components by 20 executing a corresponding one or more software agents, comprising the steps of:

receiving information for at least one of the packets;

computing an expected return for delivery of 25 said at least one packet from said information; and

directing the delivery of said at least one packet to optimize said expected return.

The present invention includes a method for routing packets of data through a network of a plurality of components comprising the steps of:

30 defining at least one algorithm having one or more parameters for routing the data;

defining at least one global performance measure of said at least one algorithm;

5 executing said algorithm for a plurality of different values of said one or more parameters to generate a corresponding plurality of values for said global performance measure;

constructing a fitness landscape from said values of said parameters and said corresponding values of said global performance measure; and

10 optimizing over said fitness landscape to generate optimal values for said at least one parameter.

Brief Description of Drawings

FIG. 1 provides a flow diagram describing the operation 15 of software agents that direct the delivery of packets of data by controlling corresponding components in a communication network.

FIG. 2 provides a flow diagram for determining optimal values of parameters of methods performing routing control and system control.

20

25

30

Detailed Description of the Preferred Embodiment

The present invention consists of installing an independent software agent at one or more routers. In the preferred embodiment, the independent software agents are installed in some or all of the routers at any level in a hierarchy of networks and subnetworks. Each software agent updates the routing information (as a non-limiting example, routing tables) in the memory of its associated router, and shares connectivity and load information with other software agents. The software agent may either run on the same processor as its associated router or on a different processor.

Each agent acts autonomously to optimize the value of some function combining its own performance index, and that of some (zero or more) selected neighbors (not necessarily immediate topological neighbors) as explained more fully below. The performance index is based on one of the following:

- (a) its "earnings" from transmitting packets; or
- (b) a local measure of communication performance such as combining indices of load on adjacent links and expected delivery times of packets passing through its router.

Agents learn to optimize their performance index using reinforcement learning. An exemplary reinforcement learning technique is Q-learning.

Without limitation, the following embodiments of the present invention are described in the illustrative context of a solution that installs software agents at the routers of a communication network. However, it will be apparent to persons of ordinary skill in the art that the present invention also applies to the use of software agents to

control other components of the communication network. For example, software agents could control one or more directional or non-directional communication links.

FIG. 1 provides a flow diagram 100 describing the operation of software agents that direct the delivery of 5 packets of data by controlling corresponding components in a communication network. In step 102, the software agent receives information on a packet of data from other software agents. Next, in step 104, the software agent computes an expected return for delivering the packet of data using the 10 information. Next, in step 106, the software agent controls the routing of the data through its corresponding component to optimize the expected return. In step 108, the software agent transmits information to other software agents so that they can similarly control their corresponding components to 15 optimize their expected return.

Integration With existing technology

As a non-limiting example, the present invention integrates with existing standards surrounding the Open Shortest Path 20 First (OSPF) routing standard (RFC-2328) as follows:

Routing tables for OSPF-compatible routers : Preferably, the agents will not make routing decisions for each and every communication request. For example, the software agents will 25 not make routing decisions for each packet that is to be routed towards some destination. Instead, the agents will modify the routing information that the routing software or hardware uses to make decisions about communication requests. Preferably, the routing information is stored in routing tables. Thus, the agent may take a significant amount of 30 time to perform a single action such as changing one entry in a routing table. Further, this single action may

subsequently affect decisions made by the router for an indefinite period of time.

5 **Hash-based load division** : As a non-limiting example, in packet-switched networks it is usually desirable to route all packets from the same source destined for the same destination along the same route. This scheme is used to prevent out-of-order arrival of packets. This scheme can be accomplished in OSPF-compatible routers by partitioning packets for the same destination host or subnet into classes based on a hash function of the source and destination host network addresses. The classes are contiguous regions of the range hash function and the borders of these regions are defined by the routing tables.

10 The hash value could also be a function of other packet header parameters such as a reward value and quality of service specifications as defined in detail below.

15 **Opaque Link State Advertisements** : Agents must be able to communicate information about local topology and load to other agents. Preferably, this information is in the form of bids for the delivery of packets. Alternatively, this information may be directly encoded. The communication of this information takes priority over regular data traffic in the network in order to ensure its timely arrival at nodes where it is needed. As a non-limiting example, this information could be packaged in Opaque Link State Advertisements packets (RFC--

20 2370).

25

30

5

Hierarchical network structure : Networks may be structured hierarchically such that the internal structure of subnets are only visible from within the network. It will be apparent to persons of ordinary skill in the art that the present invention applies to all schemes that can be used in hierarchical networks with the modification that some of the entries in the routing tables cover groups of destinations. Similarly, some of the bids are for groups of destinations.

10

Agent performance indices

15

Agents receive immediate feedback about their performance. This feedback is called a reward. However, in the reinforcement learning framework of the present invention, an agent does not merely act to optimize its immediate reward. Instead, it acts to optimize its return. In the preferred embodiment, the return includes an expected future reward that is discounted to present value. As mentioned earlier, reward is based on "earnings" in a communication market in one of the preferred embodiments called the market-based reward framework. In another preferred embodiment called the local performance reward framework, the reward is based on an index of local communication performance.

25

Market-based reward framework

30

In the market-based reward framework, each packet contains a contract to pay some amount of a "cash" equivalent to the router that delivers it to its final destination. The contracted amount is paid in full only if the packet reaches its final destination within a constraint such as a pre-

specified quality of service constraint. Preferably, a portion of the contracted amount is paid at the destination if the packet arrives outside of the specified quality of service. This portion is determined as a function of the received quality of service. Preferably, less cash is released for packets that arrive with excessively long latency (for interactive connections). Likewise, less cash is released for packets that arrive out-of-order or at widely varying intervals (for audio or video streams). At the final destination node of a packet, market-arbiter software 5 calculates the cash reward earned by the delivering software agent and the amount owed by the originating application. These rewards and bills are accumulated over time and sent out at a low frequency so as to impose only a negligible communication load on the network.

10 When reinforcement learning is used to adjust the behavior of agents, instantaneous rewards are based on the actual cash profit of the agent and optionally, the cash profit of neighboring agents (not necessarily topological neighbors) over some short past time period. Optionally, in order to prevent agents from charging arbitrary prices in 15 monopoly situations, excess profit can be removed (taxed) from those agents whose long-term discounted expected reward exceeds a predefined target.

20

Each agent communicates "bids" that specify how much it 25 will pay for packets having a particular destination, a particular specified quality of service, and a specified maximum rate to other agents. Preferably, each agent communicates the "bids" to its topologically neighboring agents. Bids may also have an expiration time. Optionally, the bids are represented by a function. Non-limiting 30 function examples include a margin, a rate, a minimum contract value, and a minimum delivery time. For example, an

agent at node B may specify that it will pay the value less 3 units for up to 800 packets per second destined for node F having a value of at least 15 units and a remaining allowable delay of 120ms. Bids stand until they expire or until the 5 node where a bid is held receives a message canceling and/or replacing the bid. Optionally, other quality of service parameters corresponding to the quality of service requirements of packets are included in the bids. For example, a higher price may be paid for packets that arrive in sequence. Bids may also specify a route. When bids 10 specify a route, agent may not sell a packet against a bid that would result in the packet returning to the same router. For example, if B submits a bid to A to deliver packets to E via the path CDAF, then A may not sell to B packets destined for E.

15 Packets that are received by a node (either from an application program at the node, or from another node) that do not conform to the parameters of an existing bid (e.g., insufficient contract value or too many in a given time period) do not require payment. Instead, these packets are owned by the agent at the node and may be sold.

20 Optionally, in addition to the agent software, nodes also execute market-arbiter software. The market-arbiter software keeps track of bids and updates and allocates payment for packets in accordance with the previously 25 discussed market rules. Optionally, bids specify "preference surfaces" that give propensities to buy or sell as probabilistic functions of qualify of service, delay, and other features. Preference surfaces were defined in co-pending patent application number 09/345,441, titled, "An Adaptive and Reliable System and Method for Operations 30 Management" and filed on July 1, 1999, the contents of which are herein incorporated by reference. Preferably, the

market-arbiter software matches preference surfaces of bidders and sellers to optimize a total "utility" for a group of packets and routers.

5 Preferably, agents make decisions based on sources of information. The decisions include:

the determination of bids and bid updates to submit to other software agents, and

the modification of the routing tables to direct packet flow so as to optimize the expected return on the routed packets.

10 The sources of information include:

bids received from other agents,

measured flows of packets through the associated router of the agent, and

15 the expected return at the router and at neighboring routers (that are not necessarily neighbors in the topographical sense).

The execution of the software agents using these market rules lead to the following network behavior:

- Agents will pay more for packets nearer the destination.

20 The agent in the destination node receives the contract value in the packet when it delivers the packet to the destination application. Consequently, it will be willing to pay a high price (near the contract value) for such packets. The agent in next-to-last node will be willing to pay a slightly lower price, and so on.

25 Packets far from their destination will be purchased for relatively little.

- It will generally cost more to send packets further.

30 Since the agent at each node along a route takes its own margin (e.g., buys packets for 8 units, and sells them for 10 units), it will cost more to send packets further. Preferably, the margins charged by agents

reflect actual establishment and/or operating costs for particular communication links.

- *Different levels of service may be provided. An agent*
5 *may maintain different bids for different levels of service. Higher levels of service such as a faster delivery time will cost more. A packet that is sent out with sufficient contract value to cover a higher level of service but that does not arrive at its destination within the specified quality of service parameter will*
10 *only be worth a reduced value to the router making the final delivery. In this situation, the originating application will be charged only the reduced value.*
- *Application programs at nodes will know how much it*
15 *costs to send a packet to a particular destination. The bids lodged at a node specify how much it costs to send a packet to a particular destination. Once the packet is in transit, even if routing costs change, intermediate nodes are still motivated to forward packets as explained further in the next paragraph.*
- *Packets are always worth sending. Even if an agent is*
20 *caught in a crunch, it is still worthwhile for the agent to sell the packets at a loss. For example, suppose an agent receives 500 packets at a price of 7 units,*
25 *expecting to be able to sell them for 9 units. Suppose further that the bid drops to 3 units before the agent can sell them. Even in this situation, the agent will sell the packets at a loss because if it retains these packets, it receives no reward at all from them as their contract value is not realized until they reach their*
30 *destination.*

- Agents will have to make predictions about future packet flow. Since decisions cannot be made about individual packets but only about bids and routing table entries, earnings will depend on the flow of packets and may fluctuate. Preferably, agents make predictions about future packet flow in order to set routing table entries so as to maximize expected return. For example, an agent may set routing table entries to forward most of the received packets to a neighbor who pays well for them (but not too many, since it will not receive a reward for the ones sold above a predetermined rate as explained in the preceding monopoly discussion).
5
- Agents will be motivated to keep bids up-to-date and high. If an agent charges too large a margin (ie., its bids are too low), it will loose business to
10 competitors, and consequently will receive a lower return. If an agent lets its bids get out-of-date and too high, it will receive a lower or negative return on packets that it forwards. Hence, agents will be
15 motivated to keep bids high (i.e. margins low) and up-to-date.
20
- Earnings at nodes can help guide decisions about short- and long-term resource allocation. If margins at nodes are designed to accurately reflect costs of
25 communication, then market theory indicates that prices charged by agents will accurately reflect benefits of allocating additional resources (barring monopoly situations). Thus, prices charged by agents can be used as a guide for allocating short-term or long-term
30 resources such as a temporary connection or a leased line.

Local-performance reward framework

An alternative to the market-based reward scheme is a scheme where local rewards are based on unbiased estimates of 5 packet delivery times. Preferably, packet delivery times are estimated in a decentralized fashion by plugging reported link loads into models of network performance. The immediate reward for an agent at a node is the inverse of an increasing function of the aggregate estimated packet delivery times. Optionally, the immediate reward also incorporates other 10 indices of quality of service. In the local performance reward framework, agents modify routing tables in an attempt to reduce the estimated delivery times or improve other aspects of quality of service.

15 Locally-cooperative local reinforcement learning

Having all agents attempt to optimize their local figures of merit will not always result in the discovery of the globally optimum configuration as explained in "At Home in the Universe" by Stuart Kauffman, Oxford University Press, 20 Chapter 11 in the context of an NK fitness landscape, the contents of which are herein incorporated by reference. This result occurs because actions taken by one agent affects its state and possibly changes the context of the reward for its neighboring agents.

25 Accordingly, in the preferred embodiment the present invention utilizes combinations of the following three semi-local strategies:

30 **patches** In this technique, agents are partitioned into disjoint subsets called patches. The patches may or may not be topologically contiguous. Within a patch, the actions of agents are coordinated to

maximize the aggregate figure of merit for the entire patch. The size and location of patches are parameters for this strategy.

5 p A neighborhood is defined for a node such that when a decision is made there, figures of merit at the current node and at a proportion p of neighboring nodes are taken into account. A neighborhood need not consist of the immediate topological neighbors of the node.

10 tau Only a fraction (called tau) of the agents make decisions that change the portions of their state that affect the reward of other agents at the same time.

15 FIG. 2 provides a flow diagram 200 for determining optimal values of parameters of methods performing routing control and system control. In step 210, the present invention defines a global performance measure for the network. In step 220, the present invention defines an optimization algorithm having at least one parameter.

20 Exemplary parameters include the size and location of patches, the neighborhood, p where the figures of merit are considered in making a decision and the fraction, tau, of the agents that change portions of their state that affect the reward of other agents. In step 230, the method 200 constructs a landscape representation for values of the parameters and their associated global performance measure. In step 240, the method optimizes over the landscape to produce optimal values for the parameters.

25 In the preferred embodiment, the present invention uses either patches or p or both to define a modified reward and hence, a return, for an agent in the network routing problem. As explained earlier, the figure of merit for an agent is

either its earnings in the market-based framework or its local measure of performance in the local performance framework. Optionally, the present invention uses the tau strategy either alone, or in conjunction with p and "patches" 5 to limit the opportunities agents have for making decisions that affect the return of other agents. For example, the reward for an agent is the aggregate earnings for a region of agents (a patch) and the bids and routing tables for only a fraction tau of agents change at the same time.

Preferably, the parameters for these strategies (the fraction 10 p , the fraction tau and the number and membership of patches) are global in nature. In other words, the values of these parameters are the same for all agents. Alternatively, the values of the parameters may vary among the agents.

Preferably, the present invention sets these parameters 15 as follows:

First, a global performance measure is defined. Preferably, the global performance measure is a combination of the average delivery time and the achieved network bandwidth. Second, the algorithm has an outer loop that varies these parameters in order to maximize the global 20 performance measure in accordance with techniques for searching landscapes as described in the co-pending international patent application titled, "A System and Method for the Analysis and Prediction of Economic Markets", filed December 22, 1999 at the U.S. receiving office, the contents 25 of which are herein incorporated by reference.

Preferably, each value of the global parameters governing p , patches, tau, and reinforcement learning features, defines a point in the global parameter space. With respect to this point, the bandwidth-agent system of the present invention achieves a given global fitness. The 30 distribution of global fitness values over the global parameter space constitutes a "fitness landscape" for the

entire bandwidth-agent system. Such landscapes typically have many peaks of high fitness, and statistical features such as correlation lengths and other features as described in co-pending international patent application number 5 PCT/US99/19916, titled, "A Method for Optimal Search on a Technology Landscape", the contents of which are herein incorporated by reference. In the preferred embodiment, these features are used to optimize an evolutionary search in the global parameter space to achieve values of p , patches, tau, and the internal parameters of the reinforcement 10 learning algorithm that optimize the learning performance of the bandwidth-agent system in a stationary environment with respect to load and other use factor distribution. Preferably, the same search procedures are also used to persistently tune the global parameters of the bandwidth-agent system in a non-stationary environment with respect to 15 load and other use factor distributions.

By tuning of the global parameters to optimize learning, the present invention is "self calibrating". In other words, the invention includes an outer loop in its learning procedure to optimize learning itself, where co-evolutionary 20 learning is in turn controlled by combinations of p , patches, and tau, plus features of the reinforcement learning algorithm. The inclusion of features of fitness landscapes aids optimal search in this outer loop for global parameter values that themselves optimize learning by the bandwidth-agent system in stationary and non-stationary environments. 25

Use of p , tau, or patches aids adaptive search on rugged landscapes because, each by itself, causes the evolving system to ignore some of the constraints some of the time. Judicious balancing of ignoring some of the constraints some of the time with search over the landscape optimizes the 30 balance between "exploitation" and "exploration". In particular, without the capacity to ignore some of the

constraints some of the time, adaptive systems tend to become trapped on local, very sub-optimal peaks. The capacity to ignore some of the constraints some of the time allows the total adapting system to escape badly sub-optimal peaks on the fitness landscape and thereby, enables further searching. In the preferred embodiment, the present invention tunes p , τ , or patches either alone or in conjunction with one another to find the proper balance between stubborn exploitation hill climbing and wider exploration search.

The optimal character of either τ alone or patches alone, is such that the total adaptive system is poised slightly in the ordered regime, near a phase transition between order and chaos. See e.g. "At Home in the Universe" by Kauffman, Chapters 1, 4, 5 and 11, the contents of which are herein incorporated by reference and "The Origins of Order", Stuart Kauffman, Oxford University Press, 1993, Chapters 5 and 6, the contents of which are herein incorporated by reference. For the p parameter alone, the optimal value of p is not associated with a phase transition.

Without limitation, the embodiments of the present invention are described in the illustrative context of a solution using τ , p , and patches. However, it will be apparent to persons of ordinary skill in the art that other techniques that ignore some of the constraints some of the time could be used to embody the aspect of the present invention which includes defining an algorithm having one or more parameters, defining a global performance measure, constructing a landscape representation for values of the parameters and their associated global performance value, and optimizing over the landscape to determine optimal values for the parameters. Other exemplary techniques that ignore some of the constraints some of the time include simulated annealing, or optimization at a fixed temperature. In general, the present invention employs the union of any of these means to ignore some of the constraints some of the

time together with reinforcement learning to achieve good problem optimization.

Further, there are local characteristics in the adapting system itself that can be used to test locally that the system is optimizing well. In particular, with patches alone and tau alone, the optimal values of these parameters for adaptation are associated with a power law distribution of small and large avalanches of changes in the system as changes introduced at one point to improve the system unleash a cascade of changes at nearby points in the system. The present invention includes the use of local diagnostics such as a power law distribution of avalanches of change, which are measured either in terms of the size of the avalanches, or in terms of the duration of persistent changes at any single site in the network.

The present invention's use of any combination of the above strategies, together with reinforcement learning in any of its versions, give it an advantage over prior art routing methods because these strategies address many problems that could arise including the following:

20 - slow convergence to optimal routing patterns,
- oscillation of network load, and
- locally beneficial but globally harmful routing patterns.

Without limitation, the embodiments of the present invention have been described in the illustrative context of a method for routing data through a communication network. However, it is apparent to persons of ordinary skill in the art that other contexts could be used to embody the aspect of the present invention which includes defining an algorithm having one or more parameters, defining a global performance measure, constructing a landscape representation for values 30 of the parameters and their associated global performance

value, and optimizing over the landscape to determine optimal values for the parameters.

For example, the present invention could be used for operations management as explained in co-pending U.S. patent application No. 09/345,441, titled, "An Adaptive and Reliable System and Method for Operations management" and filed on July 1, 1999, the contents of which are herein incorporated by reference. That patent describes a model of an enterprise in its competitive environment, based on technology graphs that support a nodes and flow model of an organization, plus a management structure. The present invention, using agents 10 to represent objects and operations in the enterprise model, coupled to reinforcement learning, p , patches and tau, is used advantageously to create a model of a learning organization that learns how to adapt well in its local environment. By use of the outer loop described above, good 15 global parameter values for p , patches, tau, and the reinforcement learning algorithm are discovered. In turn, these values are used to help create homologous action patterns in the real organization. For example, the homologous action patters can be created by tuning the 20 partitioning the organization into patches, by tuning how decisions at one point in the real organization are taken with respect to a prospective benefit of a fraction p of the other points in the organization affected by the first point, and by tuning what fraction, tau, of points in the organization should try operational and other experiments to 25 improve performance.

In addition, the distribution of contract values and rewards in the reinforcement algorithm can be used to help find good incentive structures to mediate behavior by human agents in the real organization to achieve the overall 30 adaptive and agile performance of the real organization.

5 In addition to the use of the invention to find good global parameters to instantiate in the real organization, the same invention can be used to find good global parameter values to utilize in the model of the organization itself to use that model as a decision support tool, teaching tool, etc.

Further, the present invention is also applicable to portfolio management, risk management, scheduling and routing problems, logistic problems, supply chain problems and other practical problems characterized by many interacting factors.

10 While the above invention has been described with reference to certain preferred embodiments, the scope of the present invention is not limited to these embodiments. One skill in the art may find variations of these preferred embodiments which, nevertheless, fall within the spirit of the present invention, whose scope is defined by the claims set forth below.

15

20

25

30

Claims

1. A method for routing packets of data through a
5 network of a plurality of components comprising the steps of:
controlling one or more of said components by
executing a corresponding one or more software agents,
comprising the steps of:
receiving information for at least one of the
packets;
- 10 computing an expected return for delivery of
said at least one packet from said information; and
directing the delivery of said at least one
packet to optimize said expected return.
- 15 2. A method as in claim 1 wherein said
information for said at least one packet comprises a
destination.
- 20 3. A method as in claim 2 wherein said
information for said at least one packet further comprises a
contract to pay a specified reward to said one or more
software agents that delivers said at least one packet to
said destination.
- 25 4. A method as in claim 3 wherein said
information of said at least one packet further comprises a
specified quality of service.
- 30 5. A method as in claim 4 wherein said specified
reward varies with a delivered quality of service in
comparison with said specified quality of service.

6. A method as in claim 4 wherein said information for said at least one packet comprises at least one bid specifying a price that said one or more software agent will pay for said at least one packet having said 5 destination and said quality of service.

7. A method as in claim 4 wherein said quality of service comprises a latency for said at least one packet.

8. A method as in claim 4 wherein said quality of 10 service comprises a specified order for delivery of said at least one packet.

9. A method as in claim 1 wherein said information for said at least one packet comprises at least one bid specifying a price that said one or more software 15 agent will pay for said at least one packet.

10. A method as in claim 9 wherein said at least one bid further comprises an expiration time.

20 11. A method as in claim 9 wherein said at least one bid further comprises a margin.

12. A method as in claim 9 wherein said at least one bid further comprises a minimum value.

25 13. A method as in claim 9 wherein said at least one bid further comprises a minimum delivery time.

14. A method as in claim 9 wherein said at least one bid further comprises a specified route.

30

15. A method as in claim 9 wherein said at least one bid is a satisfaction profile defining a satisfaction of trading said at least one packet as a probability density function of at least one parameter.

5

16. A method as in claim 15 wherein said at least one parameter of said probability density function comprises a quality of service.

17. A method as in claim 1 wherein said expected 10 return for delivery of said at least one packet is an expected reward discounted to present value.

18. A method as in claim 1 wherein said expected return for delivery of said at least one packet step varies 15 inversely with an estimated delivery time for said at least one packet.

19. A method as in claim 18 wherein said controlling one or more components step further comprises the step of transmitting delivery loads to others of said one or 20 more software agents for determining said estimated delivery time for said at least one packet.

20. A method as in claim 1 wherein said one or more software agents control one or more legal entities of 25 the network.

21. A method as in claim 1 wherein said one or more software agents control one or more communication links of the network.

30

22. A method as in claim 1 wherein said controlling one or more of said components step further

comprises the step of partitioning said one or more software agents into one or more patches.

23. A method as in claim 22 wherein said directing the delivery of said at least one packet step comprises the 5 step of optimizing said expected return of said patch.

24. A method as in claim 1 wherein said computing an expected return step comprises the step of:

selecting a portion p of said one or more software 10 agents; and

computing said expected return of said selected portion p of said one or more software agents.

25. A method as in claim 24 wherein said delivery 15 of said at least one packet is directed to optimize said expected return of said selected portion p of said one or more software agents.

26. A method as in claim 1 wherein said controlling one or more of said components step further 20 comprises the step of transmitting said information from said one or more software agents to others of said software agents.

27. A method as in claim 26 wherein said 25 transmitted information comprises at least one bid specifying a price that said one or more software agents will pay for said at least one packet.

28. A method as in claim 26 wherein said transmitted information comprise delivery loads.

30

29. A method as in claim 26 wherein only a fraction, tau, of said one or more software agents transmit said information at the same time.

5 30. A method for routing packets of data through a network of components comprising the steps of:

defining at least one algorithm having one or more parameters for routing the data;

defining at least one global performance measure of said at least one algorithm;

10 executing said algorithm for a plurality of different values of said one or more parameters to generate a corresponding plurality of values for said global performance measure;

15 constructing a fitness landscape from said values of said parameters and said corresponding values of said global performance measure; and

optimizing over said fitness landscape to generate optimal values for said at least one parameter.

31. A method as in claim 30 wherein said defining 20 an algorithm step comprises the steps of:

controlling one or more of said components by executing a corresponding one or more software agents comprising the steps of:

25 communicating information for at least one of the packets among said one or more software agents;

computing an expected return for delivery of said at least one packet from said information; and

directing the delivery of said at least one packet to optimize said expected return.

32. A method as in claim 31 wherein said at least one parameter comprises a proportion p of said one or more software agents.

5 33. A method as in claim 32 wherein said computing an expected return step comprises the step of:

computing said expected return of said proportion p of said one or more software agents.

10 34. A method as in claim 31 wherein said at least one parameter comprises a size of one or more patches of said one or more software agents and a location of said patches.

15 35. A method as in claim 34 wherein said directing the delivery of said at least one packet step comprises the step of:

optimizing said expected return of said patch.

20 36. A method as in claim 31 wherein said at least one parameter comprises a fraction, τ , of said one or more software agents.

37. A method as in claim 36 wherein only said fraction, τ , of said software agents communicate information for said at least one packet at the same time.

25 38. A method for performing operations management in an environment of entities comprising the steps of:

representing at least one of the entities with at least one corresponding model having a plurality of parameters;

30 defining at least one global performance measure of said model;

executing said model for a plurality of different values of said at least one parameters to generate a corresponding plurality of values for said global performance measure;

5 constructing a fitness landscape from said values of said parameters and said corresponding values of said global performance measure; and

 optimizing over said fitness landscape to generate optimal values for said at least one parameter.

10 39. A method as in claim 38 wherein said representing at least one of the entities with at least one corresponding model having a plurality of parameters step comprises the steps of:

15 representing a plurality of decision making units within the entities with a corresponding plurality of decision making agents; and

 representing a plurality of communication links among the decision making units with a corresponding plurality of connections among said plurality of decision making agents.

20 40. A method as in claim 39 further comprising the steps of:

 communicating information among said decision making agents;

25 computing an expected return at said decision making agents from said information; and

 making at least one decision at said decision making agent to optimize said expected return.

30 41. A method as in claim 40 wherein said at least one parameter comprises a proportion p of said decision making agents.

42. A method as in claim 41 wherein said computing an expected return step comprises the step of:

computing said expected return of said proportion p of said decision making agents.

5

43. A method as in claim 40 wherein said at least one parameter comprises a size of one or more patches of said decision making agents and a location of said patches.

10 44. A method as in claim 43 wherein said making at least one decision step comprises the step of:
optimizing said expected return of said patch.

15 45. A method as in claim 40 wherein said at least one parameter comprises a fraction, τ , of said decision making agents.

46. A method as in claim 45 wherein only said fraction, τ , of said decision making agents communicate information at the same time.

20 47. Computer executable software code stored on a computer readable medium, the code for routing packets of data through a network of a plurality of components, the code comprising:

25 code to control one or more of said components by executing a corresponding one or more software agents, comprising:

code to receive information for at least one of the packets;

30 code to compute an expected return for delivery of said at least one packet from said information; and

code to direct the delivery of said at least one packet to optimize said expected return.

48 50. A programmed component for routing packets of data through a network comprising at least one memory having at least one region storing computer executable program code and at least one processor for executing the program code stored in said memory, wherein the program code comprises:
5 code to control one or more of said components by executing a corresponding one or more software agents,
10 comprising:

code to receive information for at least one of the packets;

code to compute an expected return for delivery of said at least one packet from said information;
15 and

code to direct the delivery of said at least one packet to optimize said expected return.

49. Computer executable software code stored on a computer readable medium, the code for routing packets of data through a network of a plurality of components, the code comprising:
20

code to define at least one algorithm having one or more parameters for routing the data;

25 code to define at least one global performance measure of said at least one algorithm;

code to execute said algorithm for a plurality of different values of said one or more parameters to generate a corresponding plurality of values for said global performance measure;

30 code to construct a fitness landscape from said values of said parameters and said corresponding values of said global performance measure; and

code to optimize over said fitness landscape to generate optimal values for said at least one parameter.

5 50. A programmed component for routing packets of data through a network comprising at least one memory having at least one region storing computer executable program code and at least one processor for executing the program code stored in said memory, wherein the program code comprises:

code to define at least one algorithm having one or more parameters for routing the data;

10 10 code to define at least one global performance measure of said at least one algorithm;

code to execute said algorithm for a plurality of different values of said one or more parameters to generate a corresponding plurality of values for said global performance measure;

15 15 code to construct a fitness landscape from said values of said parameters and said corresponding values of said global performance measure; and

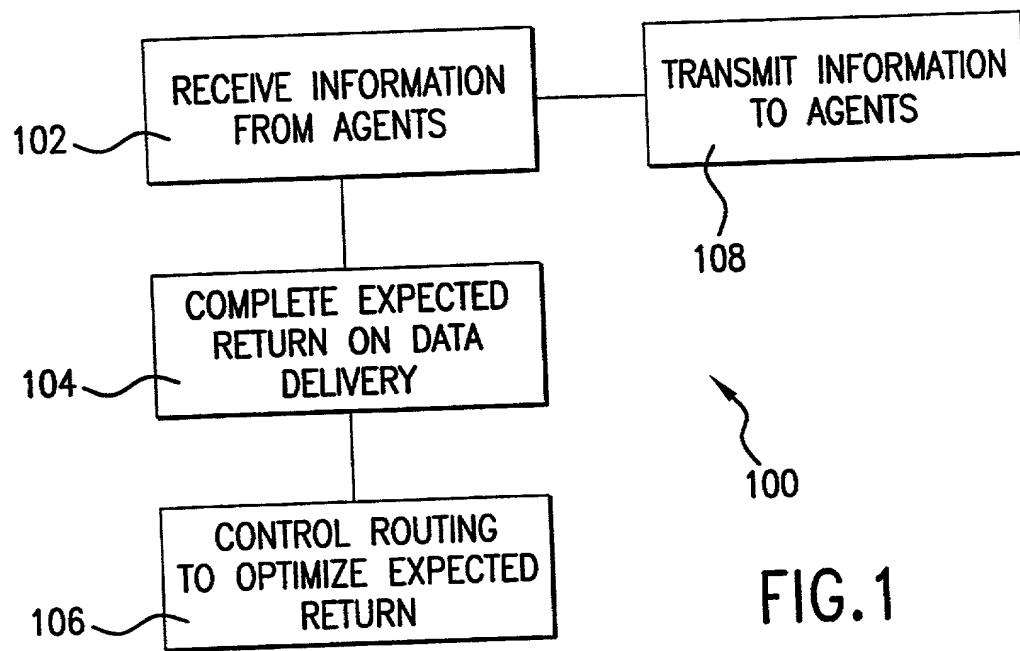
code to optimize over said fitness landscape to generate optimal values for said at least one parameter.

20

25

30

1/2



2/2

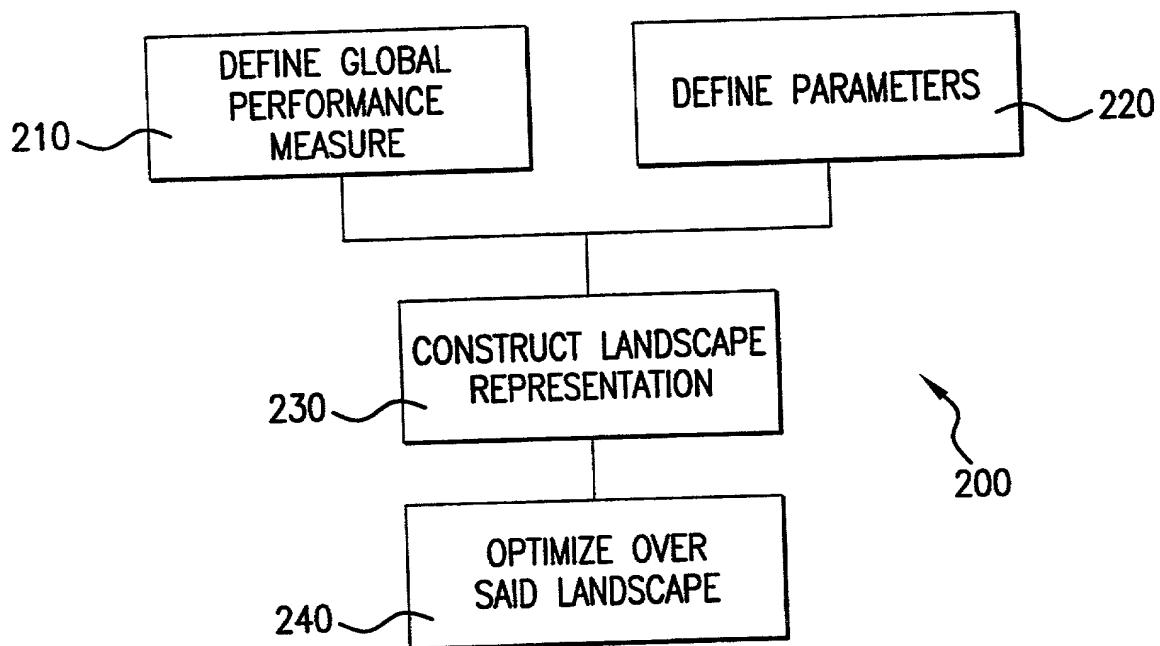


FIG.2

DECLARATION FOR NON-PROVISIONAL PATENT APPLICATION*

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below at 201 et seq. beneath my name.

I believe I am the original, first and sole inventor if only one name is listed at 201 below, or an original, first, and joint inventor if plural names are listed at 201 et seq. below, of the subject matter which is claimed and for which a patent is sought on the invention outlined

A METHOD AND SYSTEM FOR ROUTING CONTROL IN COMMUNICATION NETWORKS AND FOR SYSTEM CONTROL

and for which a patent application:

- is attached hereto and includes amendment(s) filed on 11/20/2000 (if applicable)
- was filed in the United States on 11/20/2000 as Application No. 09/688,860 (or docket number 4394-040-222) with amendment(s) filed on 11/20/2000 (if applicable)
- was filed as PCT International Application No. PCT/US00/02011 on January 28, 2000 and was amended under PCT Article 19 on 11/20/2000 (if applicable)

I hereby state that I have reviewed and understand the contents of the above identified application, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119(a)-(d) of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

EARLIEST FOREIGN APPLICATION(S), IF ANY, FILED PRIOR TO THE FILING DATE OF THE APPLICATION			
APPLICATION NUMBER	COUNTRY	DATE OF FILING (day, month, year)	PRIORITY CLAIMED
			YES <input type="checkbox"/> NO <input type="checkbox"/>
			YES <input type="checkbox"/> NO <input type="checkbox"/>
			YES <input type="checkbox"/> NO <input type="checkbox"/>

I hereby claim the benefit under Title 35, United States Code, §119(e) of any United States provisional application(s) listed below:

PROVISIONAL APPLICATION NUMBER	FILING DATE
60/117,663	January 28, 1999

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code §112, I acknowledge the duty to disclose information known to me which is material to patentability as defined in Title 37, Code of Federal Regulations, §1.56 which became available between the filing date of the prior application and the national or PCT International filing date of this application:

NON-PROVISIONAL APPLICATION SERIAL NO.	FILING DATE	STATUS		
		PATENTED	PENDING	ABANDONED
PCT/US00/02011	January 28, 2000		✓	

* for use only when the application is assigned to a company, partnership or other organization.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.

2 0 1	FULL NAME OF INVENTOR	LAST NAME KAUFFMAN	FIRST NAME Stuart	MIDDLE NAME A.
	RESIDENCE & CITIZENSHIP	CITY Santa Fe	STATE OR FOREIGN COUNTRY NM	COUNTRY OF CITIZENSHIP United States
	POST OFFICE ADDRESS	STREET 1811 S. Camino Cruz Blanca	CITY Santa Fe	STATE OR COUNTRY NM
SIGNATURE OF INVENTOR 201			DATE 7/27/01	
2 0 2	FULL NAME OF INVENTOR	LAST NAME PLATE	FIRST NAME Tony	MIDDLE NAME A.
	RESIDENCE & CITIZENSHIP	CITY Santa Fe	STATE OR FOREIGN COUNTRY NM	COUNTRY OF CITIZENSHIP Australia
	POST OFFICE ADDRESS	STREET 124 Elena Street	CITY Santa Fe	STATE OR COUNTRY NM
SIGNATURE OF INVENTOR 202			DATE 7/27/2001	
2 0 3	FULL NAME OF INVENTOR	LAST NAME	FIRST NAME	MIDDLE NAME
	RESIDENCE & CITIZENSHIP	CITY	STATE OR FOREIGN COUNTRY	COUNTRY OF CITIZENSHIP
	POST OFFICE ADDRESS	STREET	CITY	STATE OR COUNTRY
SIGNATURE OF INVENTOR 203			DATE	
2 0 4	FULL NAME OF INVENTOR	LAST NAME	FIRST NAME	MIDDLE NAME
	RESIDENCE & CITIZENSHIP	CITY	STATE OR FOREIGN COUNTRY	COUNTRY OF CITIZENSHIP
	POST OFFICE ADDRESS	STREET	CITY	STATE OR COUNTRY
SIGNATURE OF INVENTOR 204			DATE	
2 0 5	FULL NAME OF INVENTOR	LAST NAME	FIRST NAME	MIDDLE NAME
	RESIDENCE & CITIZENSHIP	CITY	STATE OR FOREIGN COUNTRY	COUNTRY OF CITIZENSHIP
	POST OFFICE ADDRESS	STREET	CITY	STATE OR COUNTRY
SIGNATURE OF INVENTOR 205			DATE	